

Adaptation and Learning in. Robotics and Automation

K.S. Narendra

S.T. Venkataraman

Center for Systems Science¹

Dept. of Electrical Engineering

Yale University

P.O.Box 208267, New Haven, CT 06520-8267

Abstract

The paper attempts to describe the reciprocal impact that adaptive control, learning theory, robotics, and automation have had on each other in the past and the influence they may have on each other in the future.

1 Introduction

To the nontechnical person, robotics and automation conjure up visions of human-like performance by machines, increased productivity, and improved quality of products. To the technically oriented person they imply electromechanical device technologies, data acquisition through sophisticated sensors, computation using powerful microprocessors, large memories, and software. To the automatic control theorist they represent broad and challenging areas of application whose demands for faster and accurate controllers are having a profound impact on the development of control theory itself. These different views of robotics and automation are accurate in their respective contexts and it is safe to say that advances in both of them will have a major impact on life in the future. In this paper, we confine our attention to an assessment of the mutual impact that adaptive control and learning theory, and robotics and automation have had on each other in the past, and attempt to predict possible influences they may have on each other in the future.

? Definitions and concepts

To establish a common framework within which questions that arise can be discussed, we briefly consider in this section the definitions and concepts the four fields have given rise to. The underlying principle in all four cases is control.

Control: By the control of a process we mean qualitatively the ability to direct, alter, or improve its behavior, and a control system is one in which some physical quantities are maintained more or less around prescribed values (regulation) or time-varying functions (tracking). The most fundamental concept in control theory and its distinctive hallmark is feedback. It underlies the whole technology of automatic control. Besides feedback, the key concepts of control theory are sensitivity, dynamic stability, and optimality.

¹S.T. Venkataraman is currently a Visiting Fellow at the Center, on assignment from the NASA Jet Propulsion Laboratory, where he is a member of technical staff in the Automation and Control Section.

The performance of any control system, however complex it may be, is judged by its speed of response and by its accuracy. Hence, what is needed in any control situation is a controller which assures stability and robustness of the control system, while being fast and accurate.

Adaptation: In control theory, adaptation is defined as the process by which a system monitors its own performance, when some aspect of the system is either unknown or changing with time, and adjusts its control parameters to improve its performance.

Learning: is a highly controversial term which means many things to different people. In engineering, learning may be characterized as the gradual change in behavior pattern under a series of exposures to the same situation. Different types of learning have been studied in the engineering literature. These include supervised learning (or learning with a teacher), reinforcement learning, and unsupervised learning. In supervised learning, the "correct" answer is provided by a teacher (in the form of an error signal). In reinforcement learning, the learning agent merely receives a signal called the reward for the action that it chooses. In unsupervised learning even the reinforcement signal is not present.

Robotics: The term originated from the Czech word "robota" and came into common use after Karel Capek's play "Rossum's Universal Robots" in 1920. In the years that followed, popular imagination suggested diverse definitions and tasks to be performed by robots of humanoid appearance. Research and development were directed towards synthetically creating biological behavior, although most real-life robots are merely electromechanical devices in the form of multi-jointed manipulators or mobile machines.

Automation: This term was first used in 1936 by D.S. Harder to describe "automatic handling of parts between progressive production processes." Attempting to make the best possible use of available resources including man, material, money, and machine, it is concerned with the linking of machine tools with automatic materials transfer and handling equipment to achieve self-regulation.

For the sake of convenience, in the remainder of this section as well as in the following sections, we confine our discussion to only two of the four terms, i.e., adaptation and robotics. Even though stated differently, learning paradigms are closely related to adaptive control paradigms and hence we use a single term to denote both of them. Only in section 6 do we explicitly refer to learning as distinct from adaptation. Between the robots of academic research and automation exists a wide spectrum of industrial robots which manifest to different degrees the properties of both. Hence, once again, we use a single term, robot, generically to denote the entire class of such systems.

Both adaptation and robotics deal with control problems under uncertainty. While control theorists view robotics as specific applications of control theory, the perception of roboticists is quite different. The latter view control theory in general (and adaptation and learning in particular) as merely tools in their arsenal to achieve their objectives.

There is also an essential difference between the research carried out in these areas. A adaptive control theory is methodology-driven, and the principal objective over the past twenty years has been to obtain conditions for global stability (or convergence in the case of learning). While stability is undoubtedly an important requirement, one cannot help wondering whether this objective has been pursued without an equal amount of time being expended on performance.

In contrast to the above, robotics, to a large extent, is problem driven where the emphasis is on achieving high levels of system performance through electromechanical design, sensing, and

computation rather than on theory. In fact, many of the important contributions in robotics thus far have resulted from good design involving sound control principles and common sense, rather than any sophisticated analysis.

Further, unlike adaptive control, there is no single set of problems which, when resolved, would let the field of robotics move forward. The inherent system dynamics of a manipulator is distinctly different from that of a hopping robot or a small Mars rover. However, as the operational speeds of these different robots are increased and the specifications on their performance become stringent, stability questions are bound to become paramount, making the kind of mathematical analysis typical of control theory indispensable for good robot design.

Adaptive control theory has found wide application in many fields, where linearized equations describe the behavior of the system sufficiently accurately. Surprisingly, even though highly nonlinear manipulators satisfy conditions required for applying the algorithms developed in the adaptive literature. The reasons for this are described in section 3.

When many of the problems in robotics, which are at present being investigated experimentally, reach a stage when mathematical analysis becomes essential, it is the belief of the authors that they will be described by nonlinear differential equations with considerable amounts of uncertainty. In such cases, adaptive control theory as it presently exists cannot be directly applied. In the second half of the paper, a general methodology which has proved very successful in simulation studies is proposed to address such problems.

Finally, due to space limitations, a list of references is not included at the end of the paper. However, names of contributors to important ideas are given throughout the paper to facilitate the search by the interested reader for relevant published literature.

3 Adaptive Control and Robotics

a) Adaptive Control: The adaptive control problem of linear time-invariant systems may be stated as follows: A plant P is represented by a linear time-invariant system

$$\begin{aligned}\dot{x}_p &= A_p x_p + b_p u \\ y &= h_p^T x_p\end{aligned}\tag{1}$$

where u is the input, y is the output, and $x_p \in \mathcal{R}^n$ is the state of the plant. The elements of h_p , A_p and b_p are assumed, to be unknown. Using only input and output information and without, using differentiators, the objective is to determine a control function which stabilizes the overall system and results in an output $y(t)$ which tracks the desired output $y_m(t)$ asymptotically, i.e., $\lim_{t \rightarrow \infty} |y(t) - y_m(t)| = 0$. The desired output $y_m(t)$ is the output of a reference model with transfer function $W_m(s)$, whose input is any bounded piece-wise continuous function $r(t)$. It took several years to determine the conditions on the plant, as well as $W_m(s)$, which are sufficient to determine a solution.

Because of space limitations, we present in this section only the highlights of the results obtained in the past two decades, in the area of adaptive control.

(i) If $e = x_p - x_m$ and ϕ denotes the parameter error, the evolution of e can be described by an equation of the form $\dot{e} = f_1(e, \phi)$. The objective is to determine an "adaptive law" $\dot{\phi} = f_2(e, t)$

where f_2 is independent of ϕ , so that the overall system is stable or asymptotically stable. A Lyapunov function $V(e, \phi) > 0$ with $\dot{V} \leq 0$ is used for the purpose.

(ii) **The** first problem to be solved in the late 1960s was for the case when (1) is a scalar system. The adaptive law has the form $\dot{\phi} = -e w$ where w is an accessible signal.

(iii) Shortly thereafter the problem where the system is of n^{th} order was resolved, provided all states of the plant are accessible. The adaptive laws contain products of known signals and linear combinations of the elements of the error vector e .

(iv) Since $\dot{V}(e, \phi)$ can never be negative definite, methods similar to LaSalle's invariance principle have to be used to prove that the error tends to zero.

(v) Around this time the importance of persistency of excitation of the reference input for parameter convergence was realized. Conditions for persistent excitation were derived.

(vi) In the early 1970s interest shifted to the case when all state variables of the plant are not accessible. The stability of "adaptive observers", which simultaneously estimate the state and parameters of the system in (1) from input-output data was derived in 1973.

(vii) Around 1978, the stability of the adaptive control problem was solved for the case when the relative degree n^* of the plant is one. The adaptive laws have the same simple form as in the earlier cases.

(viii) In 1980, the same problem, for $n^* > 1$ was solved using an augmented error in place of the actual output error.

(ix) Since simulation studies revealed that adaptive systems were not robust under perturbations, a great part of the 1980s was spent in developing robust methods.

(x) In the 1980s, the methods were extended to multivariable systems.

(xi) In the 1980s and early 1990s, the results were extended to nonlinear systems, linear in parameters, when all state variables are accessible.

Work is currently in progress in [ix], (x) and (xi) as well as in relaxing the conditions on the controlled process to permit stable adaptive control.

In summary, the research of the past two decades has provided conditions under which linear and nonlinear dynamical systems can be adaptively controlled in a stable fashion. On the negative side, the methods only apply to systems in which the unknown parameters occur linearly. There is very little control over the transient response, and adaptation is practically successful only when parameters vary slowly over time.

(b) Adaptive Robot Control: Adaptive control theory has found most application in the control of robot manipulators due to certain characteristics inherent in its dynamics. The dynamic behavior of a manipulator is given by the equation

$$M(\theta)\ddot{\theta} + C(\theta, \dot{\theta})\dot{\theta} + G(\theta) = U \quad (2)$$

where θ is a vector of joint angles, $M(\theta)$ is the inertia matrix, $C(\theta, \dot{\theta})$ contains the coriolis and centrifugal terms, $G(\theta)$ is the gravity vector, and U is a vector of input torques. Associated with

u are motor dynamics, friction, backlash and saturation.

(i) An important property enjoyed by the system given in equation (2) is that the map from U to $\dot{\theta}$ is passive.

(ii) If p denotes a q dimensional vector of parameters (masses, first and second moments of the links of the robot), equation (2) can be expressed as

$$Y(\theta, \dot{\theta}, \ddot{\theta})p = U \quad (3)$$

where $Y(\theta, \dot{\theta}, \ddot{\theta})$ is a known matrix, if $\theta, \dot{\theta}, \ddot{\theta}$ are all accessible. Since equation (3) is linear in p , adaptive laws for estimating p can be readily generated, using well known methods in adaptive control.

(iii) The most general result in adaptive robot control involves a control law of the form $U = Y_d \hat{p} - \hat{K}_p e - \hat{K}_v \dot{e}$, with parameter update rules for \hat{p} using Y_d and linear terms in e, \dot{e} , and tuning laws for \hat{K}_p and \hat{K}_v involving linear terms in e, \dot{e} , where $Y_d = Y(\theta_d, \dot{\theta}_d, \ddot{\theta}_d)$, and $e = \theta - \theta_d$.

This result has had significant impact on the control of robot manipulators, and is applicable to any robotics problem involving rigid, multiple bodies. However, the basic problem addressed is a relatively simple one from the point of view of adaptive control theory (all state variables accessible), and uses results that were available in the 1970s. Therefore, application of the full arsenal of results developed in adaptive control can be expected to move the field of robot control even further. For example, adaptive control theory has developed a rich set of tools using estimates of states, derived from observers. Recent work in robotics has started to apply such results where only θ is known, and $\dot{\theta}$ is estimated using an observer. At the same time practical problems that arise in robotics due to backlash and saturation of the control input have motivated adaptive control theorists to investigate their effects on the stability and robustness of adaptive systems.

4 Robotics in Complex Environments

In the preceding section we described how results from adaptive control theory have found direct application in robot manipulator control. However, as mentioned earlier, many of the truly interesting advances in robotics have been conceptual in nature and have been in the area of practical design rather than theory. The method suggested by Brooks, Miller, etc., that uses combinations of simple control strategies to achieve seemingly intelligent behavior, the design concepts of Whitney and Raibert to stabilize the system by proper interaction of mechanical components with the active parts of the system, the ideas of Hogan (Salisbury and Mason) based on impedance (stiffness and compliance) in matching between robot and environment, the hopping robots of Raibert, the juggling robots of Koditschek, the brachiating robots of Fukuda and the devil sticking robots of Atkeson which exemplify intermittent control, are some of the truly ingenious examples in the field. These, and efforts by others to expand their scope, are constantly throwing challenges to the control theorist to suggest improved methods of control. As time progresses the same robots will be called upon to operate under increasingly stringent conditions. The equations that describe the dynamics of the overall system are already nonlinear, and quite often all state variables of the system are not accessible. In some situations sensor data is not available at a central location and in others control has to be based on patterns of data rather than the data itself. In addition to the above, there are also multiple goals, difficult to quantify constraints, and multiple time-scales to contend with.

All these, particularly when substantial amounts of information essential for mathematical study are not available, will thrust upon us questions of adaptive and learning control which cannot be answered using existing theory.

The difficulties enumerated above can be broadly classified under three headings: (i) computational complexity, (ii) uncertainty and (iii) nonlinearity. Successful autonomous operation in such complex situations will require a carefully designed interconnection of diverse information processing capabilities to achieve fast, accurate, stable and robust performance.

5 Neural Networks

The term neural network has come to mean any architecture that has massively parallel interconnection of simple "neural processors". From a system theoretic point of view, artificial neural networks are practically implementable convenient parametrizations of nonlinear maps from one finite dimensional space to another. Their parallel distributed architecture, their ability to approximate nonlinear maps arbitrarily closely, and the availability of algorithms to train their parameters using input-output data obtained while a system is in operation, make them ideally suited to cope with all three types of difficulties mentioned in the previous section.

(a) Neural Networks for Nonlinear Control: During the past five years it has been systematically demonstrated that the complex problems of nonlinear adaptive control can be addressed using neural networks. This involves a judicious combination of the results from nonlinear control theory with the concepts and structures provided by linear adaptive control and the approximating capabilities of neural networks.

In the first stage of the above investigations it was established that exact input-output representations of finite dimensional nonlinear systems in the neighborhood of the equilibrium states of the form (with $y^*(k)$ as the desired output)

$$\begin{aligned} y(k+d) = f[y(k), y(k-1), \dots, y(k-n+1), \\ y^*(k+d), u(k-1), \dots, u(k-n+1)] \end{aligned} \quad (4)$$

are possible under certain conditions. For such systems, controllers described by equations of the form

$$\begin{aligned} u(k) = g[y(k), y(k-1), \dots, y(k-n+1), \\ y^*(k+d), u(k-1), \dots, u(k-n+1)] \end{aligned} \quad (5)$$

can also be shown to exist. If f and g in equations (4) and (5) are unknown, the problem is one of nonlinear adaptive control. The first author and his colleagues have shown that f and g can be approximated by neural networks N_f and N_g using input-output data. The method has been extended to multivariable systems, systems with external disturbances, and has also been tested on systems with time-varying parameters. In all cases, the responses of the nonlinear identifiers and controllers were at least as good as those of linear controllers and in most cases surpassed them significantly.

(b) Neural Control for Robotics: The effectiveness of neural networks for both pattern recognition and for mimicking rule-based expert systems is currently well known. From the previous section it

is also clear that they are unusually effective for the identification and control of complex nonlinear systems in the presence of uncertainty. In view of this versatility, it is believed that neural networks hold great promise as building blocks for the wide variety of behaviors encountered in complex robotic systems.

6 A New Methodology for Control

In this section we describe a new methodology for controlling complex dynamical systems which has been developed at Yale in recent years. The objective of the approach is to achieve faster and more accurate control than can be achieved using conventional adaptive control, even while assuring the robustness of the system. The approach is based on the use of multiple models, switching, and adaptation. The multiple models themselves are based on neural networks and are realized through a process of learning. In this section we provide a brief introduction to the basic concepts involved.

(a) Multiple Models: The use of an identification model for indirect adaptive control is well known. When plant parameters vary abruptly as in the case of system faults, or sensor or actuator failures, adaptation using a single model may not be adequate. In fact, extensive simulation studies of both linear and nonlinear systems have shown that with large parametric errors, the response of the system may be practically unacceptable. In such cases, what information is relevant, and how it can produce the desired control action have to be determined rapidly and accurately. This is achieved using multiple models. Assuming that N identification models $I_i (i=1, 2, \dots, N)$ are used in parallel with the given plant the objective is to determine which among them is closest (according to some criterion) to the plant at any given instant and to use a corresponding controller C_i to control the plant.

(b) Neural Networks, and Multiple Models: As described earlier when the process to be controlled is nonlinear and/or contains substantial amounts of uncertainty, neural networks are ideally suited for identification and control. In such a case nonlinear identification models I_i are chosen corresponding to region S_i in the parameter space, where the plant can lie. Assuming that at any instant a fault occurs and the plant switches from S_i to S_j it is detected by the model I_j whose error function is a minimum at that instant. Adaptation takes place within the set S_j from the model I_j .

(c) Multiple Models and Learning: The determination of the regions S_j and the location of the models $I_j \in S_j$, have to be based on the prior information concerning the process and has to be carried out over long periods of time. As long as the number of specific types of situations in which the system will be called upon to act is finite and can be learned as the system is in operation, this approach can be very effective. The creation, modification, and pruning of models, the acquisition of their sensitivity, and characteristics, and the determination of the domains S_j , raise theoretical questions for which answers are not known at present. Work is proceeding in all these areas under the direction of the first author.

7 Adaptation and Learning in Robotics and Automation

As described in section 4, the frontiers of robotics and automation are constantly expanding bringing in their wake new problems in adaptation and learning. The methodology described in sections 5

and 6 is applicable to large classes of such problems and the real interaction between adaptation and robotics will be mainly in these settings. Even though we are still very much in the initial stages, we list below a few of the areas where we believe that such interaction is imminent.

(a) Robots Operating in Multiple Environments: Robots, that are required to carry a wide range of payloads and/or come in contact with environments which have a large range of inertial, stiffness and damping coefficients, can be considered to operate in multiple environments. The method described in section 6 are directly applicable to such systems. Work done by the first author and his colleague (Ciliz and Narendra) has indicated the potential for significant improvement.

(b) Gait Control: It is now well known that biological systems perform locomotion using "gaits" which are used to stabilize their body motions. Hopping, cat thing, juggling, and brachiating robots may be considered to have their own characteristic gaits. As the terrains in which the systems operate change, switching of gaits and tuning of the parameters of the gait model may become necessary. This fits very well with the multiple model paradigm.

(c) Rover Self-Localization: One of the most challenging problems in navigation is the accurate estimation of the position of a lightweight rover from its internal sensors. What is required in such a case is a dynamic model that relates the input motor torques with the position and heading of the rover. Since, such a model will invariably be nonlinear in the presence of significant interactions with the terrain, a neural network is ideal for its realization. Multiple models may be needed if the rover is to operate in very different, terrain settings.

(d) Vision and Control: The basic idea behind "visual servoing"⁷ is to kinematically transform the control problem to the camera space, and track reference trajectories prescribed in camera coordinate system. As such, this problem is similar to performing output control of a nonlinear system. The methods prescribed in sections 5 and 6 would enable this problem to be restated using only the inputs (joint torques) and outputs (robot motion in camera coordinate system). Using neural networks, the input-output model can be approximated sufficiently closely, and in the event several cameras are used, multiple neural network models may be used for visual servoing.

(e) Intermittent Control: In juggling, hopping, brachiation, etc., control signals are required along certain degrees of freedom. only intermittently. In addition, the nature of control actions depends strongly upon the constraints imposed by the bodies that are in flight. In the event, such robots are required to operate in a wide range of environments, distinctly different constraints may arise during operation depending upon the specifics of the problem. For example, constraints during free juggling, as opposed to when the balls make contact with some surfaces during flight, may be very different. Should the designer have knowledge of the various operating conditions, models could be developed for each situation and used within the framework of multiple models, switching and tuning.

The above discussion merely provides a brief glimpse of the possible areas where adaptation and robotics may interact in the future. Many problems in multiple operator teleoperation, cellular robotics, behavior control, etc. can also be addressed using a similar approach.

In conclusion, the methodology proposed in this paper based on learning, multiple models, and adaptation is deeply rooted in control theory and addresses important questions in the control of complex systems. Since its implementation admits a very traditional mathematical model, the resulting dynamical system can be studied using analytical tools of systems theory. The authors

firmly believe that it is ideally suited to cope with many of the problems that will arise in robotics and automation in the future,

8 Acknowledgements

The work described in this paper was performed jointly by the NASA Jet Propulsion Laboratory, Caltech, under the sponsorship of National Aeronautics and Space Administration, Office of Aeronautics and Exploration Technology, and by Yale University under the sponsorship of the National Science Foundation. S.T. Venkataraman was supported by a Grant from the the Office of Naval Research during this period.